

Amendments to the Claims:

This listing of claims will replace all prior versions, and listings, of claims in the application.

Listing of Claims:

1 (Previously presented): A method implemented in a computer, the method comprising:

generating an array update operation based on a query to update a relational database;

wherein said array update operation specifies a plurality of row-identifier and value pairs to update multiple rows in a table of said relational database;

repeatedly finding, and storing in a structure, a block-identifier of a block that contains a row of data identified by a row-identifier in at least a group of row-identifier and value pairs in said plurality, by use of a database index prior to retrieval of the block;

wherein said structure is located in main memory of said computer;

wherein each value comprises data to be updated in said row identified by said row-identifier;

performing a single access operation without context switching, to retrieve from a storage device and store in a cache, a number of blocks of data of said table, said blocks being identified by a corresponding number of block-identifiers in the structure;

wherein several of said blocks are non-contiguous in said storage device; and

repeatedly updating, in blocks in the cache, each row identified in the group of row-identifier and value pairs, using a corresponding value in the row-identifier and value pairs.

2 (Previously presented): The method of claim 1 further comprising:

sorting the block-identifiers, prior to retrieval of the blocks by performing the single access operation.

3 (Previously presented): The method of claim 2 wherein:

the sorting is performed subsequent to storage of the block-identifiers in the structure.

4 (Previously presented): The method of claim 1 further comprising:

subsequent to said finding and prior to said storing, checking if the block-identifier has a duplicate already stored in the structure and if so then not storing the block-identifier in the structure.

5 (Previously presented): The method of claim 1 further comprising, prior to updating:

repeating said finding of block-identifiers for all row-identifiers in the group of row-identifier and value pairs.

6 (Previously presented): The method of claim 1 wherein:

the database index is a hash index and the table is organized in a hash cluster;
and

during said finding, a single directory is used to obtain the block-identifier.

7 (Previously presented): The method of claim 1 wherein:

the database index is a B-tree index.

8 (Previously presented): The method of claim 1 wherein:

said structure comprises an array; and

the array has a number of entries identical to the number of blocks that can be held in the cache.

9 (Currently amended): The method of claim 1 further comprising:

writing a plurality of logs, at least one log for each row identified in the group of row-identifier and value pairs and ~~performing a write operation from said cache to said storage device when space is needed in said cache.~~

10 (Currently amended): The method of claim ~~[[9]]~~ 1 further comprising, ~~during said write operation:~~

unpinning each block after updating all rows in said each block; and

flushing an unpinned block to disk said storage device only when another block needs space in the cache occupied by the unpinned block.

11 (Previously presented): The method of claim 1 wherein:

a plurality of file offsets are provided to the single access operation, one file offset for each block in the group;

wherein each file offset is an offset in a file from where reading of data is to begin.

Claim 12 (canceled).

13 (Previously presented): A non-volatile media in which are stored instructions to perform a method comprising:

generating an array update operation based on a query to update a relational database;

wherein said array update operation specifies a plurality of row-identifier and value pairs to update multiple rows in a table of said relational database;

repeatedly finding, and storing in a structure, a block-identifier of a block that contains a row identified by a row-identifier in at least a group of row-identifier and value pairs in said plurality, by use of a database index of a relational database;

performing a vector read operation without context switching during said performing, to retrieve from a storage device and store in a cache, a number of blocks, said blocks being identified by block-identifiers in the structure; and

repeatedly updating, in blocks in the cache, each row identified in the group of row-identifier and value pairs, using a corresponding value in the row-identifier and value pairs;

wherein several of said blocks are non-contiguous in said storage device.

14 (Previously presented): The non-volatile media of claim 13 being further encoded with said structure storing the block-identifiers.

15 (Currently Amended): A computer comprising a processor and a memory coupled to the processor, the memory being encoded with instructions to:

automatically generate an array update operation based on a query to update a relational database;

automatically use a database index to look up a block-identifier of a block that contains a row identified by an identifier in a plurality of identifier and value pairs to be used to perform said array update operation on a table in said relational database;

automatically store the block-identifier in a structure in memory;

automatically repeat instructions to said automatically use and said automatically store, for all identifiers in at least a group of identifier and value pairs in said plurality;

automatically perform a vector read, to retrieve from a disk and store in a cache, each block in a group of blocks identified by block-identifiers stored in said structure, wherein the group of blocks are all stored in the cache during execution of a single function call;

wherein several of said blocks in the group are non-contiguous in said disk;

automatically modify a row in a block stored in the cache, using a value in the plurality of identifier and value pairs; and

automatically repeat instructions to said automatically modify, with each row identified in the group of identifier and value pairs.

16 (Currently Amended): An apparatus comprising a database, the apparatus comprising:

means for generating an array update operation based on a query to update the database;

wherein said array update operation specifies a plurality of row-identifier and value pairs to update multiple rows in a table of the database;

means for using a database index to look up a block-identifier of a block that contains the row identified by an identifier in the plurality of identifier and value pairs;

means for storing the block-identifier in a structure in memory;

means for repeating (using the database index to look up and storing the block-identifier), for all identifiers in at least a group of identifier and value pairs;

means for performing a vector read without context switching, to retrieve from a disk and store in a cache, each block in a group of blocks identified by block-identifiers stored in said structure, wherein the group of blocks are all stored in the cache during execution of a single function call;

wherein several of said blocks in the group are non-contiguous in said disk;

means for modifying a row in a block stored in the cache, using a value in the plurality of identifier and value pairs; and

means for repeating said modifying with each row identified in the group of identifier and value pairs.

17 (Currently Amended): A method implemented in a computer, the method comprising:

generating an array update operation based on a query to update a database;

wherein said array update operation specifies a plurality of row-identifier and value pairs to update multiple rows in a table of said database;

finding a block-identifier of a block that contains the row identified by a row-identifier in a row-identifier and value pair in said plurality, by use of a database index in said database;

storing the block-identifier in a structure in memory;

repeating (finding the block-identifier and storing the block-identifier), for all row-identifiers in at least a group of row-identifier and value pairs in said plurality;

performing a vector read operation without context switching during said performing, to retrieve from a storage device and store in a cache, each block in a group of blocks identified by block-identifiers stored in said structure, wherein the group of blocks are all stored in the cache during execution of a single function call;

wherein several of said blocks in the group are non-contiguous in said storage device;

updating the row in the block in the cache, using the value in the row-identifier and value pair; and

repeating said updating with each row identified in the group of row-identifier and value pairs.

18 (Previously presented): The non-volatile media of claim 13 being comprised in at least one of a disk, a chip and a cartridge.

19 (Previously presented): The method of claim 2 wherein:

the blocks are sorted during said sorting based on adjacency such that during performance of said single access operation, block-identifiers of blocks physically adjacent to one another at a periphery of a disk in the storage device are presented at one time to the storage device and identifiers of blocks that are physically adjacent to one another and located closer to a center of the disk are presented at another time.

20 (Previously presented): The computer of claim 15 wherein:

the blocks are sorted during said single function call based on adjacency such that block-identifiers of blocks physically adjacent to one another at a periphery of said disk are presented at one time to a disk drive comprising said disk and identifiers of blocks that are physically adjacent to one another and located closer to a center of said disk are presented at another time.

21 (New): The method of Claim 9 further comprising:

performing a write operation from said cache to said storage device when space is needed in said cache.